



21 October 2024

[Redacted]
[Redacted]

Dear [Redacted]

Thank you for your request made under the Official Information Act 1982 (OIA), transferred to Inland Revenue from the Office of the Minister of Revenue, on 24 September 2024. You requested the following:

I have recently learned that the Inland Revenue Department appears to have sent personal Taxpayers' information to some Social Media companies.

Why? If this is true then the Government should be locking those people (not the taxpayers) up because it is a prison offence to do so. Can you kindly investigate and excoriate those who did it. If not, then why did this information get circulated that the IRD had been negligent.

In some circumstances Inland Revenue provides information to social media platforms in order to reach particular categories of taxpayers. That information involves the use of "custom audience lists". On 12 September we paused the use of custom audience lists while we undertake an internal review of this practice (undertaken by our Chief Information Security Office). We did this due to public concern and to provide reassurance that we take the public's concerns seriously. We are also working with the Office of the Privacy Commissioner who is assessing whether there are any privacy issues with hashing. Once completed, we will be publishing the findings of the internal review on our website (ird.govt.nz).

The Privacy Act 2020 allows personal information to be used for more than one purpose. It doesn't require you to give permission every time your personal information is used or disclosed. The Act is purpose-focused not consent-focused. This means if an organisation obtains your information for one purpose, it can use it for another purpose in some circumstances.

The Tax Administration Act 1994 permits Inland Revenue to collect and use taxpayer information to carry out or support a lawful function to administer the tax system. Using social media to display tax-related messages to taxpayers (including to persons who are in debt, have due dates for payment, or telling them about entitlements) is such a function.

Privacy Impact Assessment

A privacy impact assessment helps identify whether an activity will impact on people and their information. As it may be of interest to you, I have enclosed the recent September 2024 Privacy Impact Assessment, titled *Use of Custom Audience on social media for advertising campaigns – Privacy Threshold Assessment*, for your reference (**Appendix A**).

Publishing of OIA response

We intend to publish our response to your request on Inland Revenue's website (ird.govt.nz) as this information may be of interest to other members of the public. This letter, with your personal details removed, may be published in its entirety. Publishing responses increases the availability of information to the public and is consistent with the OIA's purpose of enabling more effective participation in the making and administration of laws and policies and promoting the accountability of officials.

Thank you against for your request.

Yours sincerely



Pip Knight
Service Leader, Marketing & Communications



Use of Custom Audience on social media for advertising campaigns

Privacy Threshold Assessment

Date: Updated September 2024

Supply ID: N/A

About this Document

The purpose of this document is to demonstrate that privacy has been considered in a project or process that involves personal information. The Analysis pulls together relevant information to determine whether a full Privacy Impact Assessment (PIA) should be completed and records IR's decision of why a PIA has not been done. It will answer the following questions:

1. Does this proposal involve a new way of managing personal information?
 2. Does the proposal raise a significant privacy risk for IR?
 3. Is a full privacy impact assessment required?
-

Project Summary

1.1 Description

Inland Revenue (IR) uses social media to direct advertising to specific customer groups. This targeted approach significantly enhances IR's ability to ensure taxpayers meet their tax obligations, such as payments or that taxpayers are informed of social policy entitlements.

Social media platforms offer a service called "custom audience" to reach relevant users. The goal of using custom audience advertising is to display messages to specific taxpayers who have been identified as being in debt, have upcoming due dates, have potential entitlements, or need to update their details.

By directing advertising to those who need to see it, IR encourages taxpayer compliance, increases tax revenue collection, and informs customers of entitlements.

Examples of where IR regularly uses custom audience advertising are:

- Student loan customers who have debt owing. We separate these audiences into different segments for accuracy, such as those based overseas, those within New

Zealand who are self-employed, and those based in New Zealand earning salary and wages.

- GST customers who have returns and/or a debt due.
- Income tax debt – customers who have tax debt and would benefit from setting up an instalment arrangement to manage this debt.
- Working for Families customers who need to update their information with IR.

IR uses Meta (Facebook and Instagram), LinkedIn and Google for direct advertising. IR does not upload custom audience lists to other social media such as X (formally known as Twitter) or TikTok.

A custom audience list is a list of specific customers who Inland Revenue wants to display advertisements to.

Social media users can control which ads are displayed to them. In both Google and Facebook, users can view the advertisers whose audiences they have been included in, limit which ads are shown and can also permanently delete activity data tied to their account. LinkedIn users can choose whether they see advertising and from which businesses.

How Custom Audience works

IR uploads into its browser a list of identifiers such as names, email addresses or postal codes belonging to individuals that it wants to target with ads. This data is hashed within the browser of the IR managed device before being uploaded to the advertising platform.

The data is transmitted from an IR device to the platforms through an encrypted channel (HTTPS e.g. TLS protocol). The advertising platform will then match the hashed data with its existing hashed database to create a matched list for an advertising campaign.

For matched hashes, users are added to a Custom Audience stored within the ad manager account. If a hash does not match, it is ignored. Once the matching process completes, all hashes – both matching and non-matching – are deleted from the social media platform. IR ends up with a “custom audience” that it can target with relevant ads. The identity of matched users is not revealed to IR. This Custom Audience is stored in the ad manager account where only authorised account administrators can access it. The IR marketing team cannot see specific individuals who are contained in the Custom Audience, they just see the list name and approximate number of people that this audience contains.

No third party will have access to the information.

How hashing works

Hashing is a commonly used technique in day-to-day digital life (e.g. logging into a banking application, checking the integrity of a file you download from the internet etc). Social media platforms pre-compute the hashed values for every user so if someone uses social media, the platform will already hold their information in hashed format.

Hashing is a cryptographic security method that involves taking a piece of data - like an email address - and using math to turn it into a number (called a hash) in a consistent way. The identifier is transformed into code like a fingerprint of the original data. For example, John.doe@ird.govt.nz may come out hashed as wLKziR/6RoXDv1MDaXLH1UNUC9nIVr97jrTnL4TcxsM=.

The hash or fingerprint is difficult to reverse back to the original data. Hashing the same input data will always create the same hash, but the hashing is one way.

After the matching process, all matched and non-matched hashes are deleted from the social media platform's servers. No further processing of the hashed values is performed beyond the match process.

Prior to using Facebook's custom audience feature, IR was provided with a 2013 independent audit report by PWC. This report agreed that specified controls were in place:

- Information provided for matching in the Custom Audience product is not shared with third parties.
- Safeguards protected the security and integrity of data and guard against the accidental or unauthorised access, use, alteration, or disclosure of data within Facebook systems.
- Information was only retained for the matching process to be completed. It was then appropriately disposed of.

Google and Meta advise that information from the customer lists is not used to enhance the profiles of a user. For instance, no data or flag is added to an individual Facebook user account as a result of being added to a Custom Audience. LinkedIn says it does not link third party data to individual member's activity data. IR's CISO team is currently seeking updated assurances from the social media platforms.

Content of advertising/messages displayed

IR's advertising messages vary, are reasonably generic and not tailored to a particular customer's circumstances. For instance, a message might say "Got an NZ student loan and live overseas? Make sure you make your payment before the 30 September due date". The ads don't say something as tailored as our direct email or letter communications with a customer e.g. "Get your student loan back on track, make sure to pay here".

1.2 Legal Authority

Privacy Act 2020

The source information used to collate a custom audience list is 'personal information' for the purposes of the Privacy Act. It is also sensitive revenue information (SRI) under the Tax Administration Act 1994 (TAA).

The information provided to social media platforms is hashed before it is uploaded, so it is not personal information for the purposes of the Privacy Act as the individual is not identified.

To comply with other Privacy Act obligations, IR informs customers how it uses information. IR's [Privacy Policy](#) states that if a customer gives IR their email address or mobile phone number, we may use these to send them reminders about their tax affairs or information about our products and services. The email address is used in compiling the custom audience list.

The Policy also informs the public that IR uses information to target advertising to them indirectly via a third party:

Why you might see a certain advertisement on social media

We may also use or disclose your information to third parties to assist us to communicate or market our services to you.

To reach groups of people with information that is relevant to them while protecting their privacy, we sometimes provide hashed and fully anonymised information to social media channels when placing advertisements. In this process, your personal information is treated with the utmost integrity by us. The social media channel is not given any identifiable information. We fully comply with our obligations under the Tax Administration Act and the Privacy Act to protect your personal information.

The Privacy Act does not require an individual to consent to the use of their information.

Tax Administration Act 1994

The TAA restricts the circumstances when tax confidential information can be disclosed. Under section 18(1) of the TAA a revenue officer must keep confidential all SRI and must not disclose the information unless the disclosure is permitted under the TAA.

- **Revenue information** means information that is acquired, obtained, accessed, received by, disclosed to, or held by the Commissioner in connection with a revenue law and
- **SRI** means revenue information that identifies, or is reasonably capable of being used to identify, a person or entity, whether directly or indirectly.

The information used for advertising purposes is SRI, whether hashed or not, as it's technically possible for the social media platforms to identify a user - the hashed information is already held by the platforms. The platforms have policies and terms & conditions relating to custom audiences. Both Google and Meta do not decrypt hashed data and information from custom audience lists is not used to enhance the profiles of a user.

Because it's considered SRI, IR can only disclose the information if disclosure is a permitted disclosure that meets the requirements of sections 18D to 18J of the TAA. The permitted disclosure that will apply to a particular advertising campaign will depend on the specifics of that advertising campaign. The permitted disclosures most likely to apply can be found in section 18D(1) ("Carrying into effect revenue laws"), section 18D(2) ("Carrying out function conferred on Commissioner"), and clause 11 of Schedule 7 ("Services necessary for effective administration of revenue laws").

1.3 Personal information to be used

In the table below, describe:

- the personal information that will be collected, used and/or disclosed
- the source of the information
- the purpose of the information for your project.

Note: "Personal information" is any information about an identifiable living person. A person doesn't have to be named for the information to be identifiable.

Type of personal information	Source of information	Purpose of using the information
Meta <ul style="list-style-type: none">• First name & surname• Email address	IR customer database	To improve advertising effectiveness and deliver the

<ul style="list-style-type: none"> • City • Postal Code • Country • Date of birth 		right message to the right customers at the right time
Google <ul style="list-style-type: none"> • First name & surname • Phone number • Email address • Country • Zip Code 	IR customer database	To improve advertising effectiveness and deliver the right message to the right customers at the right time
LinkedIn – company list <ul style="list-style-type: none"> • Company name • Company email • City • Zip Code • Country LinkedIn – individual list <ul style="list-style-type: none"> • Contact first name & surname • Email address • Country 	IR customer database	To improve advertising effectiveness and deliver the right message to the right customers at the right time

1.6 Governance

Outline who has been engaged to date including sponsor or senior leaders, groups that have been consulted and approvals/endorsement to date.

Name of person or group	Business Unit	Approved, Consulted, Informed etc
Marketing & Communications	Enterprise Services	Initiator
Privacy Officer	Enterprise Design & Integrity	Endorsed
Information Security	Enterprise Design & Integrity	Endorsed
Corporate Legal	Enterprise Design & Integrity	Consulted
Tax Counsel	Tax Counsel Office	Approved

2. Privacy assessment

2.1 Areas that are risky for privacy

Some types of projects are commonly known to create privacy risks. If the project involves one or more of these risk areas, it's likely that a full Privacy Impact Assessment (PIA) will be valuable.

Use this checklist to identify and record whether your proposal raises certain privacy risks.

Does the project involve any of the following?	Y/N	If yes, explain your response
Does the initiative involve a substantial change to an existing policy, process or system?	N	No substantial change. Traditionally IR has advertised on websites that people similar to the intended audience were likely to visit. Using custom audiences is a more targeted way of displaying relevant advertising to specific taxpayers to remind them of their obligations (for instance to repay a student loan or set up payment arrangements for debt) or entitlements (for instance being eligible for Working for Families or FamilyBoost).
Is it linked to a practice or activity that is listed on a risk register?	N	
Collection	Y/N	If yes, explain your response
Will IR be collecting personal information that it doesn't currently collect?	N	
Is collecting this information necessary for IR to carry out its functions?	Y	It is necessary for IR to collect this information so it can identify and contact customers. The purpose of advertising is to perform a function of IR to encourage taxpayer compliance or notify of entitlements.
Where or who is the information being collected from?		The information is collected directly from customers.

		<p>In their profile settings, social media users can view the advertisers whose audiences they have been included in and decide whether they want to be shown ads.</p>
Storage, security, and retention	Y/N	If yes, explain your response
<p>Does the initiative change the way personal or sensitive information is stored, secured or managed?</p>	<p>N</p>	<p>The customer list (provided as CSV file) is hashed using SHA-256 within the internet browser (IR device) automatically, then the hashed information is uploaded to the platform.</p> <p>IR social media staff cannot access the plain text or hashes within the platform.</p> <p>The hashed string does not identify an individual. The platform owner (Meta, LinkedIn or Google) does not decrypt the hash, it is a one-way transform only.</p> <p>Google and Meta advise that hashed information is not decrypted, and no data or flag is added to an individual user account as a result of being added to a Custom Audience. LinkedIn says it does not link third party data to individual member’s activity data. IR’s CISO team is currently seeking updated assurances from the social media platforms.</p> <p>All platforms use SHA-256 which is compliant with NZISM standards. This is a Secure Hashing Algorithm that has been approved by GCSB.</p>
<p>Where will the information be stored?</p>		<p>The hashed data is transmitted using HTTPs (TLS Transport Layer Security), which means that it is encrypted during transmission using the same technologies that are used for online banking. The social media platforms do not retain hashed data sent to them. It is deleted soon after the campaign has run.</p>

		The custom audience list which contains the hashed data is stored temporarily by IR but no individual is identified in the list.
Who will have access to the information?		No individual outside of IR can access any personal information.
How long will the information be retained?		Meta and Google promptly delete customer lists after the match process is completed - this is regardless of whether the hashed information is matched or not. LinkedIn automatically deletes hashed customer lists within 30 days.
Does it involve transferring personal information offshore, using a third-party contractor?		Hashed information is provided to the social media platforms which is matched to hashed information already held by the platforms. The hashed information uploaded by IR is not retained or stored by the platforms and meets encryption standards set out in the NZISM.
Use, disclosure, and accuracy	Y/N	If yes, explain your response
Is the information currently held by IR?	Y	
If yes to the above question, for what purpose does IR hold the information?		The information that is used is held by IR to identify and contact taxpayers.
Will the initiative use or disclose information for a different purpose to why it was obtained?	N	IR obtains taxpayer information to identify them and contact them about their tax affairs or social policy entitlements. Sending messages (including advertising) to taxpayers to remind them of their obligations is one of the purposes IR obtains this information and is specified in IR's Privacy Policy. Methods of contacting taxpayers include directly via myIR, letters, emails or marketing campaigns.
Will IR be sharing personal or taxpayer information with another organisation?	N	

Describe the data quality – is it accurate, consistent, and complete?		The information is provided to IR by the customer so is accurate at the time it is used. If an individual has changed address, the Postal Code, Zip Code or City may not be up to date if the individual has not advised IR of the change.
What processes are in place to ensure and maintain data integrity?	N/A	
Access and identification	Y/N	If yes, explain your response
Will the information be stored on a customer or staff member’s record?	N	
Does the initiative affect how people can access information IR holds about them?	N	<p>Due to the large number of ad campaigns that IR undertakes, it is difficult to tell an individual whether they were included in a specific campaign. Manual collation would be required to review the lists to confirm whether a particular individual was included and what, if any, information hashed.</p> <p>However, Facebook users can see on their profile if IR has displayed an ad to them. Customers can always update the settings on their social media account to stop ads being displayed.</p>
Does this involve a new way of identifying individuals?	N	
Other considerations	Y/N	If yes, explain your response
Can we achieve the project's purpose using less identifiable data?	N	Minimal information is used, and hashing pseudonymises the data.
Would people be surprised by this use of their information?		<p>Following a media report by RNZ on 9 September 2024 some people are surprised by IR’s use of custom audience lists.</p> <p>It is important for IR to contact customers to remind them of their obligations. Using social media to display a message is more cost effective than contacting each</p>

		<p>customer individually to remind them to set up repayments or pay tax. However, IRs use of custom audience lists is currently under review and IR paused its use of custom audience lists on 12 September 2024.</p> <p>IR takes steps to protect the underlying customer information by using hashing and Facebook and Google advise no data or flag is added to individual user accounts as a result of being added to a Custom Audience. LinkedIn says it does not link third party data to individual member’s activity data. IR’s CISO team is seeking updated assurances from the social media platforms that the data is not added to user profiles or used in any other way.</p>
<p>If using data that customers have freely volunteered, would your project jeopardise people providing this again in the future?</p>	<p>N</p>	<p>This is information that taxpayers are required to provide to IR. It is only used to display an ad if the customer is included in a custom audience and their profile on a social media platform contains the same data that has been hashed.</p>
<p>Does the initiative involve tracking or monitoring of movements, behaviour or communications?</p>	<p>N</p>	

3. Ethical considerations

3.1 Areas that may raise ethical issues

Using and analysing data can introduce risks around the unethical use of data. IR must ensure it has ethical data practices and processes to maintain customer trust.

Does the project use ethical data practices?	Y/N	If yes, explain your response
--	-----	-------------------------------

Is the proposal likely to result in some members of a group being treated differently to one another?	Y	Only in so far as these customers have been identified as having tax obligations so IR needs to contact them to remind them of their obligations which may include setting up repayments. Some customers are displayed messages notifying them they may be eligible for social policy payments such as Working for Families or FamilyBoost.
Will the proposal have an impact on vulnerable people or those identified as disadvantaged?	N	
How are we identifying and managing bias or discrimination?	N	No bias or discrimination likely. The populations selected are those that have debt, are being reminded of a due date or may be eligible for a social policy payment.
Can you foresee any harm to individuals in using the data in the way intended?	N	IR is displaying messages to selected customers. Seeing an advertisement is not likely to cause harm to these individuals. In their profile settings, social media users can decide the way their data is used for advertising. In Facebook they can also decide whether they want to be shown ads by specific advertisers and can stop IR from showing them ads.
Does the data to be used specifically identify Māori or a Māori collective?	N	
Have you considered how the proposal contributes to the active protection of Māori interests?	N	Ethnicity data is not collected.
Use of algorithms or AI	Y/N	If yes, explain your response
If using algorithms or AI is there confidence the output is robust, and assumptions are met?	N	

Will decisions informed by an algorithm or use of AI involve human review and evaluation?	N	
Will any automated decision-making process be regularly reviewed to make sure it's still fit for purpose?	N	

4. Risk assessment

If you answered “Yes” to any of the questions above, use the table below to give a rating – either **Low (L)**, **Medium (M)**, or **High (H)** – for each of the aspects of the project set out in the first column.

For risks that you’ve identified as Medium or High, indicate (in the right-hand column) how the project plans to lessen the risk (if this is known).

Aspect of the Project	Rating	Describe any risks and how to mitigate them
<p>Level of information handling</p> <p>L – Minimal personal information will be handled</p> <p>M – A moderate amount of personal information (or information that could become personal information) will be handled</p> <p>H – A significant amount of personal information (or information that could become personal information) will be handled</p>	Low	
<p>Sensitivity of the information</p> <p>L – The information will not be sensitive (name, IRD number, or job title)</p> <p>M – The information may be considered to be sensitive (contact details, date of birth plus</p>	Medium	<p>The information may include name, date of birth, email address or phone number. No financial or tax information is included.</p> <p>Hashed information is provided to the platforms. This hashed data is then matched to hashed data in their existing database to create a custom audience list for a particular advertisement campaign.</p>

<p>name plus IRD number, financial information, biometric data)</p> <p>H – The information will be highly sensitive (health or financial details, information about high profile individuals)</p>		<p>In order to be matched, the social media platform must already hold the hashed information. It will hold this information only if the individual has provided this information to the social media platform.</p> <p>No third party has access to the information. Google, Meta and LinkedIn policies state that no information from custom audience lists is used to enhance the profiles of a user.</p>
<p>Significance of the changes</p> <p>L – Only minor change to existing functions/activities</p> <p>M – Substantial change to existing functions/activities; or a new initiative</p> <p>H – Major overhaul of existing functions/activities; or a new initiative that’s significantly different</p>	<p>Low</p>	<p>Traditionally IR has advertised on websites that people similar to the intended audience were likely to visit. Using custom audiences is a more targeted way of displaying relevant advertising to specific taxpayers.</p> <p>These taxpayers may be non-compliant and have not responded to IR contacting them directly.</p> <p>IR also uses social media to inform taxpayers of social policy entitlements such a FamilyBoost or Working for Families.</p>
<p>Interaction with others</p> <p>L – No interaction with other agencies</p> <p>M – Interaction with one or two other agencies</p> <p>H – Extensive cross-agency (that is, government) interaction or cross-sectional (non-government and government) interaction</p>	<p>Medium</p>	<p>Use of social media includes the risk there may be an inference available to platforms that the ad recipients are student loan holders or tax filers (depending on the ad message).</p> <p>The platforms have policies and terms & conditions relating to custom audiences. Both Google and Meta state they do not decrypt hashed data and information from custom audience lists is not used to enhance the profiles of a user. For instance, no data or flag is added to an individual Facebook user account as a result of being added to a Custom Audience. LinkedIn says it does not link third party data to individual member’s activity data. IR’s CISO team is currently seeking updated assurances from the social media platforms.</p> <p>Seeing an IR ad does not necessarily mean the user is non-compliant. Messages vary, are reasonably generic and IR does general advertising at certain times of the year.</p>
<p>Public impact</p>		<p>Information is used for the purpose of channelling relevant advertising relating to</p>

<p>L – Minimal impact on IR and customers</p> <p>M – Likely to have some impact on customers due to changes to the handling of personal information; or changes may raise concern or media attention</p> <p>H – High impact on customers and the public, and concerns over aspects of project; widespread media interest likely</p>	<p>Medium</p>	<p>tax compliance or social policy entitlements to a specific group, for example, student loan holders or customers due to file tax returns.</p> <p>The use of web advertising is for the purpose of carrying out IR’s function to administer or implement the tax system by facilitating taxpayer compliance, educating taxpayers on how to get it right, and maximising revenue collection. The impact on customers is that they will see messages from IR and be reminded of their obligations or informed of social policy entitlements.</p>
--	----------------------	--

5. Summary of privacy impact

The privacy impact for this project has been assessed as:	Tick
<p>Low – There is little or no personal information involved; or the use of personal information is uncontroversial; or the risk of harm eventuating is negligible; or the change is minor and something that the individuals concerned would expect; or risks are fully mitigated</p>	
<p>Medium – Some personal information is involved, but any risks can be mitigated satisfactorily</p>	<p>✘</p>
<p>High – Sensitive personal information is involved, and/or several medium to high risks have been identified. <u>You must complete a full Privacy Impact Assessment</u></p>	
<p>Inadequate information – More information and analysis is needed to fully assess the privacy impact of the project.</p>	

6. Reasons for the privacy impact rating

The privacy impact rating is **Medium** for the following reasons:

- IR has a lawful purpose to contact taxpayers and remind them of their tax obligations or inform them of social policy entitlements.
- IR can legally use taxpayer information to support a duty or function of the Commissioner of Inland Revenue.
- IR has various channels it uses to contact taxpayers including myIR, phone, letters, text messages, emails and advertising.

- IR creates custom lists of potentially relevant audiences specific to a campaign (e.g. Student Loan, GST customers, Working for Families etc.) through data provided by customers.
- Minimal information is used to identify who may be included in a custom audience list.
- IR only shares hashed information with advertising platforms. Hashing is a commonly used technique to protect data.
- Hashing happens within the IR managed devices and the data is transmitted to the platforms through an encrypted channel (HTTPS e.g. TLS protocol).
- No third party has access to the information.
- Google and Meta policies state that no information from custom audience lists is used to enhance the profiles of a user.
- The potential risk of hashing is that a malicious actor could crack hash values from a compromised data set. The dataset is usually obtained from a cyber security breach. This is mitigated by the data not being stored, using robust techniques (SHA-256) and the information that could be compromised is already available on a user's account.

7. Document sign-off

Position	Business Unit	Sign-off Date
Business Owner	Marketing & Communications, Enterprise Services	18 Sept 2024
Privacy Officer	Enterprise Design & Integrity	18 Sept 2024
Chief Information Security Office	Enterprise Design & Integrity	18 Sept 2024